

METHOD AND APPARATUS FOR COMMUNICATING DATA WITHIN MEASUREMENT TRAFFIC

BACKGROUND OF THE INVENTION

5

Field of the Invention

This invention relates to the field of networking. In particular, the invention relates to communicating data within measurement traffic.

10 10 *Description of the Related Art*

Internetworks such as the Internet provide a best-effort service and do not reserve resources for a path. Hence, performance characteristics of the path such as delay, jitter and loss can change over time due to routing changes, congestion, and lack of connectivity, and therefore it is important to be able 15 to measure them. There are several tools available to measure the performance characteristics of a path:

Ping uses ICMP packets to measure reachability and round trip delay from a source host to a remote host.

Traceroute detects common reachability problems such as routing loops 20 and network black holes by sending ICMP packets from a source host to a destination host, and by receiving ICMP responses from intermediate routers along the path between the source host and the remote host. Each intermediate router in the path decrements the TTL value stored in the header of an ICMP packet by one; when the TTL field expires (reaches the value zero) in a router, 25 the router does not forward the packet towards the destination host. Instead, it returns the ICMP to the source host responding with a Time Exceeded response. By starting with an initial TTL value of 1 and gradually incrementing the TTL field in successive ICMP packets, the source host is able to receive an ICMP 30 response from all the routers in the path. Traceroute also computes the round trip time of each ICMP packet, hence being able to determine the round trip delay between the source host and intermediate routers.

Pathchar measures congestion of a path by estimating performance characteristics of each node along a path from a source to a destination. Pathchar also leverages the ICMP protocol's Time Exceeded response to packets whose TTL has expired. By sending a series of UDP packets of various sizes to each hop, pathchar uses knowledge about earlier nodes and the round trip time distribution to this node to assess incremental bandwidth, latency, loss, and queue characteristics across the link connected to this node.

These tools are mainly used for troubleshooting purposes. A more formal attempt to measure performance characteristics of Internet paths is being developed by the IP Performance Metrics (IPPM) working group of the Internet Engineering Task Force (IETF). The IPPM working group has specified a general framework for measuring performance characteristics of a path, including specifications for clock synchronization and for size, number and inter-transmission time of measurement packets. The IPPM working group has also specified specific performance metrics for one-way delay, one-way inter-packet delay variation, and one-way loss, among others. The goal of the IPPM measurement framework is to allow service providers and other network providers to develop and operate and inter-operable measurement infrastructure, for performance and billing purposes, among other purposes.

However, even if this measurement infrastructure is in place, a way to communicate measurements and performance characteristics of measured paths to appropriate points of the network where decisions based on those performance characteristics can be made, is needed. In addition, this communication should be efficient, i.e., it should minimize the amount of bandwidth consumed.

SUMMARY OF THE INVENTION

The invention includes methods and apparatuses for communicating data within measurement traffic. Some embodiments of the invention will consist of a sender of measurement packets. Some embodiments of the invention will consist of a receiver of measurement packets. Some embodiments of the invention will consist of both a sender and a receiver of measurement packets.

In some embodiments of the invention, measurement packets will traverse one or more paths traversing at least a portion of an internetwork.

5 In some embodiments of the invention, the measurement packet will include information for a receiver of the measurement packet to compute measurements of performance characteristics of at least a portion of the path of the measurement packet. In some embodiments, the measurement packet sizes and times between measurement packets will simulate the traffic pattern of one or more applications such as -- by way of a non-limiting example -- voice and video.

10 In an embodiment that includes a receiver of measurement packets, measurement statistics may be computed that are at least partly recomputed with the arrival of each measurement packet. This computation may include at least one of a moving average, an average based on the Robbins-Moro estimator, a window-based average, and a bucket-based average.

15 In some embodiments of the invention, the measurement packets will contain data including one or more of measurement statistics, a generic communication channel, network information, and control data directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

20 In an embodiment wherein the data includes measurement statistics, the measurement statistics may include one or more of delay, jitter and loss.

25 Some embodiments of the invention will contain clocks synchronized by GPS, IRIG, NTP or NIST. Some other embodiments will use unsynchronized clocks and will adjust for clock skew and drift by performing computations on the measurement data.

In some embodiments of the invention, paths will be selected based at least in part on at least one of: one or more of the measurement statistics from the measurement packet and one or more of the computed statistics.

These and other embodiments are described further herein.

30

BRIEF DESCRIPTION OF THE FIGURES

Fig. 1 shows some possible embodiments of devices that are communicating with each other, for example sending and receiving measurement packets.

5 Fig. 2 shows one specific detailed embodiment of two devices, where each device is sending and receiving measurement packets as well as selecting a subset of paths.

10 Fig. 3 shows an embodiment with more than two devices that are sending and receiving measurement packets to obtain measurements of performance characteristics of paths and to communicate measurements statistics about those paths.

DETAILED DESCRIPTION

15

Measurement Packets

A measurement packet is a packet sent by a sender over an internetwork that includes information necessary for the receiver of the packet to compute measurements of the performance characteristics of the path the packet has 20 traversed over that internetwork. The information includes information for a receiver of the measurement packet to compute measurements of performance characteristics of at least a portion of the path of the measurement packet; and data including one or more of measurement statistics, a generic communication channel, network information, and control data directing a receiver of the 25 measurement packet to change one or more configuration parameters of the receiver.

In some embodiments of the invention, the information included in the measurement packet to compute measurements includes at least one of a timestamp of a sending time of the packet and a number to identify the packet 30 by itself and/ to identify the relative position of the measurement packet in a sequence of measurement packets,

In some embodiments of the invention, the measurement packet is implemented using the following data structure:

```
struct MeasurementHeader {  
  
    /**  
     * A generation number. This value represents when the  
     * sender began sending. This value is a standard Unix  
     * timestamp that seconds since Jan 1, 1970 UTC.  
     */  
    uint32_t mGeneration;  
10  
    /**  
     * A sequence number for the packet. This increments each  
     * time a packet is sent and rolls over when 16 bits is  
     * exceeded.  
     */  
    uint16_t mSequence;  
  
    /**  
     * The IP address the packet is sent to.  
     */  
    uint32_t mDstAddr;  
  
    /**  
     * The send timestamp for this packet.  
     */  
    uint64_t mSendTime;  
25  
};
```

The mGeneration field is used to detect when a sending process has started a new session. This field is used by the receiver to determine that a discontinuity in the stream's sequence numbers is the result of a sender restart, rather than due to large network latencies, duplicate packets or dropped packets.

The sequence number mSequence field is incremented by one each time a packet is sent. This approach allows the receiver to deduce lost and duplicate packets by identifying missing and duplicate sequence numbers.

The mSendTime field contains the time at which the packet was sent, represented as microseconds since January 1, 1970 UTC. This field is

compared to the time the packet arrived at the receiver to determine the delay between the sender and the receiver.

In some embodiments of the invention, a plurality of one or more packets are sent over a path continuously. In some embodiments of the 5 invention, the continuous stream of packet is denoted as a measurement stream. Each measurement stream is uniquely identified by the source and destination IP addresses. The sender maintains one socket descriptor for each source IP address it sends from and writes the destination IP address into the mDstAddr field. On the receiver side, the source IP address is returned by the recv() 10 system call and the destination address is retrieved from the measurement packet.

Data Included in the Measurement Packets

In measurement packets that contain sufficient space, data will be 15 included, including one or more of measurement statistics, a generic communication channel, network information, and control data directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

Some embodiments of the invention will add a single type of data to 20 each packet. Some embodiments of the invention will use a complex data, including subpackets.

Some embodiments of the invention use subpackets that include a single byte subpacket type identifier, followed by a 2-byte length field (including the length of the type and length fields) and finally including the data that is to be 25 sent. One embodiment will store all values in network byte order. Other byte orders will be apparent to those skilled in the art. The following data structure definition describes some embodiments.

```
30   class SubPacket {
        /*
         * The type identifier for this subpacket.
         */
        uint8_t mType;
```

```
/*
 * The length of this subpacket, in network byte order.
 */
uint16_t mLength;
5      };
```

One embodiment of this invention will include data describing a momentary snapshot of the measurement statistics for a given path between a sender and a receiver.

10 In some embodiments of this invention, this data will include one or more of the following information: the source and destination IP addresses that define the path, a measurement packet size for which the statistics have been calculated as well as computed measurement statistics that are at least partly responsive to delay; computed measurement statistics that are at least partly 15 responsive to jitter and computed measurement statistics that are at least partly responsive to packet loss.

In one embodiment of this invention, these statistics will be in units of microseconds expressed as 64-bit floating-point quantities and transmitted in a standard network byte order.

20 In one embodiment of this invention, the following data structure will store the computed statistics:

```
class TunnelStatsSubPacket : public SubPacket {
    /**
     * The time that this statistic snapshot was taken (in
     * microseconds since 1970).
     */
25      uint64_t mTimestamp;

    /**
     * The source IP address of the tunnel these statistics
     apply
     * to.
     */
30      uint32_t mSrcAddr;
```

```
 /**
 * The destination IP address of the tunnel these
5   statistics
 * apply to.
 */
uint32_t mDstAddr;

 /**
 * The size of measurement packet that these statistics
10  apply
 * to. A size of 0 indicates that these statistics
apply to
 * all packet sizes.
 */
15  uint16_t mPktSize;

 /**
 * The average delay in microseconds.
 */
20  double mDelay;

 /**
 * The average jitter in microseconds.
 */
25  double mJitter;

 /**
 * The percentage of packets that have been lost, in the
range
30  * 0 to 1.
 */
double mLoss;
};
```

35 Some embodiments of this invention include the time at which the statistics were computed such that those statistics are sent over multiple paths for improved reliability and to take advantage of one path having less delay than another. One embodiment at the receiving end is able to compare the

computation times of received statistics to place them in their original temporal order, regardless of their relative arrival times over the paths.

Some embodiments of this invention will send computed statistics specific to the paths that are part of the plurality of one or more paths that are 5 between the specific sender and receiver. Other embodiments will send additional computed statistics for paths that are not one of the plurality of one or more paths that are between the specific sender and receiver.

Some embodiments of this invention will include network information concerning network topology including but not limited to information retrieved 10 from routers such as in-bound or out-bound link utilization, inbound or out-bound link bandwidth and/or CPU utilization. Other network information determined from routers and other network devices will be apparent to someone skilled in the art.

Some embodiments of this invention will also include control data 15 directing a receiver of the measurement packet to change one or more configuration parameters of the receiver.

In some embodiments of the invention, the control data will instruct a receiver to alter its configuration, including but not limited to zero or more of the following examples: instructing a receiver to initiate sending a plurality of 20 one or more measurement packets, change one or more of the measurement packet sizes, inter-measurement packet transmission times and mix of packet sizes, and stop sending one or more of the plurality of measurement packets.

In some embodiments of the invention, this control information will include notification of measurement devices that have joined or left the 25 network.

In many embodiments of the invention, the measurement packets will be encrypted by the sender and decrypted by the receiver. Some of these embodiments will use IPSec.

In some embodiments of the invention, the encryption and decryption 30 will be done by an external device using IPSec.

Other encryption and decryption options will be apparent to one skilled in the art.

In some embodiments of the invention, the measurement packets will be digitally signed.

In some embodiments of the invention, a generic communication channel will be used by a sender and a receiver to communicate data between them.

5

Performance Characteristics of a Path

10

Measurements are used to compute performance characteristics of the paths traversed by the measurement packets. The measurements can either be computed from the measurement packets themselves, or extracted from the arbitrary data carried by the measurement packets. The measurements of performance characteristics include at least one or more of one-way measurements and round-trip measurements. The performance characteristics include at least one or more reachability, delay, jitter, loss, available bandwidth, and total bandwidth. Other performance characteristics will be apparent to those skilled in the art.

15

In some embodiments of the invention, delay measurements are computed as the interval of time from the moment the measurement packet is sent by the sender to the moment of time the measurement packet is received by the receiver. The sending time is carried by the packet, and it is measured by the clock the sender refers to. The receiving time is measured by a clock that the receiver refers to, which may or may not be synchronized with the sender's clock.

20

In some embodiments of the invention, the clock of the sender and the clock of the receiver are synchronized. A plurality of one or more precise clock inputs such as GPS, NTP, IRIG and NIST will be used. Some embodiments of this invention will use the same clock as an input to more than one of the plurality of one or more senders and receivers. In some embodiments of the invention, the clock of the sender and the clock of the receiver are the same.

25

In some embodiments of the invention, the clock of the sender and the clock of the receiver are not synchronized, and mechanisms based on the measurement data are used to correct the clock skew and clock drift, the

mechanisms including using minimum delay across multiple measurement samples, and using a mechanism to track the minimum delay over time.

Some embodiments of the invention will use the minimum round-trip delay between the devices to place a lower bound on clock skew.

5 Some embodiments of the invention will use the lower bound of multiple paths between the sender and receiver to further reduce the lower bound.

Some embodiments of the invention will correct for clock drift by tracking the relative clock skew between the sender and receiver over time and 10 adjusting for the slope of the drift.

In some embodiments of the invention, jitter measurements, also known as inter-measurement packet delay variations, are computed as the difference in delay on consecutive, successfully received packets.

15 In some embodiments of the invention, jitter can also be computed as the difference between the instantaneous delay of a packet, and the average delay of all the measurement packets previously received.

20 In some embodiments of the invention, loss measurements are computed by assigning a timeout value to each measurement packet that indicates the instant of time after which the measurement packet will be declared lost, if the packet has not arrived by that time. In some embodiments of the invention, the timeout value of a measurement packet can be computed with the transmission time of a previously received packet, an estimation of the inter-transmission time between measurement packet, and an estimation of the transmission delay of the measurement packet. In some embodiments of the invention, the inter- 25 transmission time can be estimated if the receiver knows about the scheduling pattern of transmission of measurement packets. In some embodiments of the invention, the transmission delay of packet can be estimated based on delay and jitter performance characteristics.

30 Performance characteristics of a path could be the measurement themselves, or statistics on those measurements. In the statistics case, a dynamic algorithm is used to update the statistics associated with a path with every new measurement obtained with the arrival of every new packet over the path.

In some embodiments of the invention, the algorithm computes statistics over the performance characteristics of the path.

5 In some embodiments of the invention, the statistics include averages, deviations, and variances. Other statistics will be apparent to those skilled in the art. In some embodiments of the invention, averages can be computed using a plurality of one or more techniques including a moving average, an average based on the Robbins-Moro estimator, a window-based average or a bucket-based average. Other techniques to compute averages will be apparent to those skilled in the art.

10 In some embodiments of the invention, the moving average is an exponentially moving average computed using a Robbins-Moro estimator. The Robbins-Moro stochastic approximation estimator finds a solution of the equation:

$$E[f(t) - x] = 0$$

15 where E is the expectation, f(t) a function and x the estimator. The general form of the solution is:

$$x(t) = x(t-1) + \alpha * [f(t) - x(t-1)] = (1 - \alpha) * x(t-1) + \alpha * f(t)$$

or, with $\alpha = (1 - \mu)$,

$$x = \mu * x + (1 - \mu) * f$$

20

μ is the weight of the estimator, and determines the amount contributed to the average by the function.. In some embodiments of the invention, μ is constant. In some embodiments of the invention, μ is a dynamic value, whose value depends on the last value of the function f according to the formula:

25 $\mu = e^{-(-f/K)}$

where K is a constant that also determines the importance of the last value of f with respect to the current value of the estimator x.

In some embodiments of the invention, average delay can be computed using an exponentially moving average as follows,

30

$$d = \mu * d + (1 - \mu) * m$$

where d is the exponentially moving average of delay, m is the last delay sample, and μ is the weight of the moving average.

5

In some embodiments of the invention, average jitter can be computed using an exponentially moving average as follows,

$$v = \mu * v + (1 - \mu) * |d - m|$$

10

where v is the exponentially moving average of jitter, $|d - m|$ is the last sample of jitter, and μ is the weight of the average.

15

In some embodiments of the invention, average jitter can be computed using an exponentially moving average as follows,

$$v = \mu * v + (1 - \mu) * |m - m'|$$

20

Where v is the exponentially moving average of jitter, $|m - m'|$ is the last sample of jitter, m is the last delay sample, m' is the previous delay sample, and μ is the weight of the average.

In some embodiments of the invention, delay and jitter averages can be combined into a single value as follows:

$$l = d + M * v$$

25

Where d is the average delay, v is the average jitter and M is a constant.

In some embodiments of the invention, average loss can be computed using an exponentially moving average as follows,

30

$$\hat{p} = \mu * \hat{p} + (1 - \mu) * p$$

where \hat{p} is the moving average of the loss, $p = \{0 \text{ if packet is received, } 1 \text{ if the packet is declared lost}\}$, and μ is the weight of the exponentially moving average.

In some embodiments of the invention, μ is determined based on the
5 notion of forgiveness against a single packet loss. The forgiveness period is the interval of time between the time the packet loss occurs and the time the average loss is forgiven. The forgiveness period can be either defined in units of time, or in number of packets if the rate of the monitoring flow is known. That is, the forgiveness period will end after n consecutive packets have been
10 received after the loss, when these packets have been transmitted at a certain rate.

The value of the exponentially moving average after receiving the n packets is needed before μ can be determined, and this value is known as the
15 forgiveness threshold. In some embodiments of the invention, the forgiveness threshold is chosen arbitrarily. In some embodiments of the invention, the
forgiveness threshold takes the value:

$$\frac{1}{2} (1 - \mu)$$

This value is half of the value of the estimator after the single loss occurs, and thus we call it the *half-life threshold*. Similarly, we also call the forgiveness period under this threshold the *half-life period*. The advantage of using a
20 forgiveness threshold greater than zero is that issues related to host-dependent floating-point representations reaching that value are avoided.

In some embodiments of the invention, μ is computed by comparing the value of the estimator after n consecutive packet arrivals since the loss with the *half-life threshold*:
25

$$\hat{p} = (1 - \mu) * \mu^n < \frac{1}{2} (1 - \mu)$$

Given that n is known because is determined by the value of the *half-life period* and the transmission rate, μ is computed as:

$$\mu = \exp((\ln \frac{1}{2}) / n)$$

In some embodiments of the invention, two thresholds are defined, an upper threshold and a lower threshold. When the value of p-hat exceeds the upper threshold, the loss is not forgiven until enough measurement packets are received consecutively so that the value of p-hat gets below the lower threshold.

5

Other mechanisms to compute μ will be apparent to those skilled in the art.

Path Description

10 In some embodiments of the invention, the path traversed by the measurement packets from the sender to the receiver is such that the path is at least partly implemented with at least one of a GRE tunnel, an IPSEC tunnel and IPonIP tunnel. Other path implementations using tunnel will be apparent for those skilled in the art.

15 In some embodiments of the invention, the path traversed by the measurement packets from the sender to the receiver is implemented with a virtual circuit, including a frame relay PVC, an ATM PVC or MPLS. Other path implementations using virtual circuits will be apparent for those skilled in the art.

20 Other path implementations will be apparent to those skilled in the art.

Internetes Description

25 In some embodiments of the invention, the internetwork is implemented by a plurality of one or more subnetworks, including a plurality of one or more VPNs, a plurality of one or more BGP autonomous systems, a plurality of one or more local area networks, a plurality of one or metropolitan area networks, and a plurality of one or more wide area networks.

In some embodiments of the invention, the internetwork is implemented by an overlay network.

30 Other internetwork implementations will be apparent to those skilled in the art.

Packet Sizes and Transmission Times

In some embodiments of the invention, the measurement packets are of varying sizes, including 64, 256, 512, 1024, 1500 bytes.

5 In some embodiments of the invention, the size of the measurement packets is specified with an external API.

In some embodiments of the invention, the measurement packets are of a fixed size.

10 In some embodiments of the invention, the measurement packet sizes and times between measurement packets simulate the traffic pattern of a plurality of one or more applications

15 In some embodiments of the invention, traffic patterns correspond to voice applications, where the packets are of small size, e.g., 30 bytes, and the inter-transmission time between consecutive packets is constant, e.g., 10 ms. These examples do not limit the possible size values and inter-transmission time values.

20 In some embodiments of the invention, traffic patterns correspond to video applications, where the packets size is the largest permitted to be transmitted by an internetwork without being fragmented, and the inter-transmission time between consecutive packets varies depending on the spatial and temporal complexity of the video content being transmitted, the compression scheme, the encoding control scheme.

25 In some embodiments of the invention, traffic patterns correspond to the plurality of applications observed in an internetwork, including at least one or more of HTTP transactions, FTP downloads, IRC communications, NNTP exchanges, streaming video sessions, VoIP sessions, videoconferencing sessions and e-commerce transactions. Other types of applications will be apparent to those skilled in the art.

30 In some embodiments of the invention, the inter-measurement packet transmission times are of varying length.

In some embodiments of the invention, the inter-measurement packet transmission times are of fixed length.

In some embodiments of the invention, the inter-measurement packet transmission times specified with an external API.

In some embodiments of the invention, the length of the inter-measurement packet transmission times is randomized according to a distribution. In some embodiments of the invention, this distribution is based at least in part on a uniform distribution. In some embodiments of the invention, 5 this distribution is based at least in part on an exponential distribution. In some embodiments of the invention, this distribution is based at least in part on a geometric distribution. Other distributions will be apparent to those skilled in the art.

In some embodiments of the invention, the length of the inter-measurement packet transmission times is provided by a table. 10

In some embodiments of the invention, the length of the inter-measurement packet transmission times is controlled by a scheduler. In some embodiments of the invention, the scheduler uses a priority queue, keyed on desired send time.

15 Other mechanisms to specify the inter-measurement packet transmission time will be apparent to those skilled in the art.

Other packet sizes and transmission times will be apparent to those skilled in the art.

20 Path Selection

It is possible that multiple alternative paths between a sender and a receiver are available through an internetwork at any given moment.

Performance characteristics of each of these paths can be used to select a subset of the paths.

25 In some embodiments of the invention, the subset of the plurality of paths is selected based at least in part on at least one of: one or more of the measurement statistics from the measurement packet and one or more of the computed statistics.

30 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the position of paths in a ranking. In some embodiments of the invention, the ranking is at least partly based on one or more of the measurement statistics included as data in the measurement packet. In some embodiments of the invention the ranking is at least partly

based on the computed statistics of the path. In some embodiments of the invention the ranking is implemented by using a comparison function to compare the paths, and by ordering the paths in a decreasing order. In some embodiments of the invention the ranking is implemented by using a comparison function to compare the paths, and by ordering the paths in an increasing order. Other ranking techniques will be apparent to those skilled in the art.

In some embodiments of the invention, the ranking is based on a single score associated to each path. In some embodiments of the invention, this score is denoted *Magic Score* (MS), and it is computed as follows:

$$MS = ML * MF$$

$$ML = d + M * \nu$$

$$MF = \delta * p\text{-hat} + I$$

where ML is the *Magic Latency*, a component of the MS obtained using delay and jitter respectively calculated with statistics; and MF is the *Magic scaling Factor* that multiplies the value of ML, and is computed based on loss statistics. M is a constant that takes several values, including 4, for example. MS can be seen as a scaled-up version of ML, and the scaling factor MF is a function of p-hat and delta, a constant. As p-hat not only reflects loss but also detects large delay spikes before they happen, p-hat can be seen as an indicator of the departure of the path from a "normal mode" operation, and thus the scaling factor is only applied when there are loss or spikes. The goal of MF is to differentiate between paths that have very similar delay characteristics, but with one having losses and the other not having them.

In some embodiments of the invention, ML is used as a delay indicator, given that jitter is accounted as an increase in delay. In contrast, MS, although a scaled version of ML, cannot be used to indicate delay, except when MF = 1 (p-hat = 0), which leads to MS = ML. That means the value of MS is useful not by itself but to compare it with the MSs of other tunnels.

In some embodiments of the invention, loss statistics can be used as a discriminator instead of a scaling factor. That is, p-hat can eliminate paths experiencing loss. Then, the remaining paths can be selected using $MS = ML$.

5 In some embodiments of the invention, the selection of a subset of paths is based on applying at least one or more thresholds to at least one of more of the statistics.

10 In some embodiments of the invention, a single threshold is used, and computed as a certain percentage of the highest score of the paths. In some embodiments of the invention, the threshold is determined by subtracting a fixed quantity to the highest score of the paths.

15 In some embodiments of the invention, the number of paths in the subset of paths is fixed. In some embodiments of the invention, this fixed number of paths N out of M paths is determined such that the probability of having loss in $(M - N)$ paths simultaneously is less than a certain threshold. In some embodiments of the invention, this probability is a binomial, with the assumption that all paths have the same probability of loss.

20 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on a probability associated with each path. In some embodiments of the invention, the probability of each path is at least partly based on one or more of the measurement statistics included as data in the measurement packet.

In some embodiments of the invention, the probabilities of each path are equal.

25 In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the cost of the path.

In some embodiments of the invention, the selection of the subset of the plurality of paths is based at least partly on the amount of bandwidth consumed over a period of time.

30 Other possibilities to compute path probabilities will be apparent to those skilled in the art.

Other mechanisms to select a subset of the paths will be apparent to those skilled in the art.